Katarzyna Bien · Ingmar Nolte · Winfried Pohlmeier

A multivariate integer count hurdle model: theory and application to exchange rate dynamics

Abstract In this paper we propose a model for the conditional multivariate density of integer count variables defined on the set $\mathbb{Z}^n$. Applying the concept of copula functions, we allow for a general form of dependence between the marginal processes, which is able to pick up the complex nonlinear dynamics of multivariate financial time series at high frequencies. We use the model to estimate the conditional bivariate density of the high frequency changes of the EUR/GBP and the EUR/USD exchange rates.

Keywords Integer count hurdle · Copula functions · Discrete multivariate distributions · Foreign exchange market

JEL Classification G10 · F30 · C30

1 Introduction

In this paper we propose a model for the multivariate conditional density of integer count variables. Our modelling framework can be used for a broad set of applications to multivariate processes where the primary characteristics of the variables are: first, their discrete domain spaces, each being the whole space $\mathbb{Z}$; and second, their contemporaneous dependence.

Katarzyna Bien
University of Konstanz, Konstanz, Germany

Ingmar Nolte
University of Konstanz, CoFE, Konstanz, Germany

Winfried Pohlmeier
University of Konstanz, CoFE, ZEW, Konstanz, Germany

Ingmar Nolte (✉)
Department of Economics, University of Konstanz,
Box D124, 78457 Konstanz, Germany
E-mail: Ingmar.Nolte@uni-konstanz.de

Although econometric modelling of univariate processes with a discrete support has been studied extensively, the multivariate counterpart is still underdeveloped. Most of the existing approaches (e.g. Kocherlakota and Kocherlakota (1992) and Johnson et al. (1997)) concentrate on the parametric modelling of multivariate discrete distributions with a non-negative domain and a non-negative contemporaneous dependence only. Alternatively, Cameron et al. (2004) exploit the concept of copula functions to derive a more flexible form of the bivariate distribution for non-negative count variables that allows for both a positive or a negative dependence between the discrete random variables. The multivariate integer count hurdle model (MICH) proposed here can be viewed as a combination of the copula approach by Cameron et al. (2004) with the Integer Count Hurdle (ICH) model of Liesenfeld et al. (2006), which allows for the dynamic specification of a univariate conditional distribution with discrete domain $\mathbb{Z}$.

Quite a number of applications of the MICH model are conceivable in many academic disciplines. Most apparent are applications to high frequent financial data, which are characterized by a set of contemporaneously correlated trade marks, many of them are discrete in nature at high or ultra high frequency. In empirical studies on financial market microstructure, characteristics of the multivariate time-varying conditional densities (moments, ranges, quantiles, etc.) are crucial. For instance, with our model we are able to derive multivariate conditional volatility or liquidity measures. As an application, we propose a model for the bivariate process of exchange rate changes sampled at the 1 min frequency. Other possible applications would be, for example, modelling joint movements of stock transaction prices or the changes of the bid and ask quotes of selected financial instruments.

The discreteness of price changes plays an important role for financial theory and applications. Huang and Stoll (1994), Crack and Ledoit (1996) and Szpiro (1998) among others show that discrete price changes imply a ray shaped pattern in the scatter plot of returns against one period lagged returns, which is referred to as the ‘compass rose’. The compass rose can be found for many financial instruments on different markets, such as futures (Lee et al. (1999)), exchange rates (Gleason et al. (2000) and Szpiro (1998)) and stocks (Crack and Ledoit (1996) and Antoniou and Vorlow (2005)).

It has several implications for the dynamics of the data generating process of asset returns which may render naively applied statistical tests such as the Brock et al. (1996) test (Krämer and Runde (1997)), random walk tests or simple autocorrelation estimates (Fang (2002)) invalid. Moreover, GARCH models estimated for such data may be misspecified (Amilon (2003)) and the assumption of a geometric Brownian Motion as the true price process can at least be questioned, which has consequences, for instance, for option pricing (Ball (1988)) and the discrimination between the market microstructure noise and the underlying price process in the realized volatility literature (Andersen et al. (1999), Oomen (2005) and Hansen and Lunde (2006)). Furthermore, Vorlow (2004) analyzes to which extent such patterns can be exploited for forecasting issues. Our approach nicely contributes to this literature since the MICH is able to pick up complex nonlinear structure such as the compass rose in a multivariate setting.

The data used in the application part of the paper are 1 min changes of the EUR/GBP and the EUR/USD midquotes. Figure 1 shows its bivariate histogram. The changes of exchange rates are discrete, since bid or ask quotes of the GBP and the USD against the EUR can jump by a multiple of a fixed tick size of 0.0001

Fig. 1 Bivariate histogram of the tick changes of the EUR/GBP and the EUR/USD exchange rates.

EUR only. The bid quotes (and the ask quotes, analogously) are aggregated to the 1 min level by taking the average of the highest and the lowest best bid within that minute, resulting in a smallest bid quote change of 0.00005 EUR, so that the smallest observable mid quote change amounts to 0.000025 EUR.

Due to the discreteness of the bivariate process, the surface of the histogram is rough, characterized by distinct peaks with the most frequent outcome (0,0) having a sample probability of 2.02%, that corresponds to the simultaneous zero movement of both exchange rates. The discrete changes of the variables are positively correlated, since the positive (negative) movements of the EUR/GBP exchange rate go along with the positive (negative) movements of the EUR/USD exchange rate more frequently.

The sequence of the paper is organized as follows. In Section 2 we describe the general framework of our multivariate modelling approach. The description of the theoretical settings customized with respect to modelling the bivariate density of exchange rate changes follows in Section 3. There, we also present the results of empirical application as well as some statistical inference. Section 4 discusses the results and concludes.

2 The general model

Let $Y_t = (Y_{1t}, \dots, Y_{nt})' \in \mathbb{Z}^n$, with $t = 1, \dots, T$, denote the multivariate process of n integer count variables and let $\mathcal{F}_{t-1}$ denote the associated filtration at time $t - 1$. Moreover, let $F(y_{1t}, \dots, y_{nt} | \mathcal{F}_{t-1})$ denote the conditional cumulative density function of Y_t and $f(y_{1t}, \dots, y_{nt} | \mathcal{F}_{t-1})$ its conditional density. Each marginal process Y_{kt} , $k = 1, \dots, n$ is assumed to follow the ICH distribution of Liesenfeld et al. (2006) and the dependency between the marginal processes is modelled with a copula function.

2.1 Copula function

The copula concept of Sklar (1959) has been extended by Patton (2001) to conditional distributions. In that framework the marginal distributions and/or the copula function can be specified conditional on $\mathfrak{F}_{t-1}$, so that the conditional multivariate distribution of Y_t can be modelled as

$$F(y_{1t}, \dots, y_{nt} | \mathcal{F}_{t-1}) = \mathbf{C}(F(y_{1t} | \mathcal{F}_{t-1}), \dots, F(y_{nt} | \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}), \quad (1)$$

where $F(y_{kt} | \mathcal{F}_{t-1})$ denotes the conditional distribution function of the k th component and $\mathbf{C}(\cdot | \mathcal{F}_{t-1})$ the conditional copula function defined on the domain $[0, 1]^n$. This approach provides a flexible tool for modelling multivariate distributions as it allows for the decomposition of the multivariate distribution into the marginal distributions, which are bound by a copula function, being solely responsible for their contemporaneous dependence.

If the marginal distribution functions are continuous, the copula function $\mathbf{C}$ is unique on its domain $[0, 1]^n$, because the random variables Y_{kt} , $k = 1, \dots, n$ are mapped through the strictly monotone increasing functions $F(y_{kt} | \mathcal{F}_{t-1})$ onto the entire set $[0, 1]^n$. The joint density function can then be derived by differentiating $\mathbf{C}$ with respect to the continuous random variables Y_{kt} , as:

$$f(y_{1t}, \dots, y_{nt} | \mathcal{F}_{t-1}) = \frac{\partial^n \mathbf{C}(F(y_{1t} | \mathcal{F}_{t-1}), \dots, F(y_{nt} | \mathcal{F}_{t-1}) | \mathcal{F}_{t-1})}{\partial y_{1t} \dots \partial y_{nt}}, \quad (2)$$

However, if the random variables Y_{kt} are discrete, $F(y_{kt} | \mathcal{F}_{t-1})$ are step functions and the copula function is uniquely defined not on $[0, 1]^n$, but on the Cartesian product of the ranges of the n marginal distribution functions, i.e. $\bigotimes_{k=1}^n \text{Range}(F_{kt})$ so that it is impossible to derive the multivariate density function using Eq. (2).

Two approaches have been proposed to overcome this problem. The first is the continuation method suggested by Stevens (1950) and Denuit and Lambert (2005), which is based upon generating artificially continued variables $Y_{1t}^*, \dots, Y_{nt}^*$ by adding independent random variables $U_{1t}, \dots, U_{nt}$ (each of them being uniformly distributed on the set $[-1, 0]$) to the discrete count variables $Y_{1t}, \dots, Y_{nt}$ and which does not change the concordance measure between the variables (Heinen and Rengifo (2003)).

The second method, on which we rely, has been proposed by Meester and MacKay (1994) and Cameron et al. (2004) and is based on finite difference approximations of the derivatives of the copula function, thus

$$f(y_{1t}, \dots, y_{nt} | \mathcal{F}_{t-1}) = \Delta_n \dots \Delta_1 \mathbf{C}(F(y_{1t} | \mathcal{F}_{t-1}), \dots, F(y_{nt} | \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}), \quad (3)$$

where Δ_k , for $k \in \{1, \dots, n\}$, denotes the k th component first order differencing operator being defined through

$$\begin{aligned} \Delta_k \mathbf{C}(F(y_{1t} | \mathcal{F}_{t-1}), \dots, F(y_{kt} | \mathcal{F}_{t-1}), \dots, F(y_{nt} | \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}) \\ = \mathbf{C}(F(y_{1t} | \mathcal{F}_{t-1}), \dots, F(y_{kt} | \mathcal{F}_{t-1}), \dots, F(y_{nt} | \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}) \\ - \mathbf{C}(F(y_{1t} | \mathcal{F}_{t-1}), \dots, F(y_{kt} - 1 | \mathcal{F}_{t-1}), \dots, F(y_{nt} | \mathcal{F}_{t-1}) | \mathcal{F}_{t-1}). \end{aligned}$$

The conditional density of Y_t can therefore be derived by specifying the cumulative distribution functions $F(y_{1t}|\mathcal{F}_{t-1}), \dots, F(y_{nt}|\mathcal{F}_{t-1})$ in Eq. (3).

2.2 Marginal Processes

The integer count hurdle (ICH) model that we propose for the modelling of the marginal processes is based on the decomposition of the process of the discrete integer valued variable into two components, i.e. a process indicating whether the integer variable is negative, equal to zero or positive (the direction process) and a process for the absolute value of the discrete variable irrespective of its sign (the size process). We present here the simplest form of the ICH model and we refer to Liesenfeld et al. (2006), reprinted in this volume, for a more elaborate presentation.

Let π_{jt}^k , $j \in \{-1, 0, 1\}$ denote the conditional probability of a negative $P(Y_{kt} < 0|\mathcal{F}_{t-1})$, a zero $P(Y_{kt} = 0|\mathcal{F}_{t-1})$ or a positive $P(Y_{kt} > 0|\mathcal{F}_{t-1})$ value of the integer variable Y_{kt} , $k = 1, \dots, n$, at time t . The conditional density of Y_{kt} is then specified as

$$\begin{aligned} f(y_{kt}|\mathcal{F}_{t-1}) \\ = \pi_{-1t}^k \mathbb{1}_{\{Y_{kt} < 0\}} \cdot \pi_{0t}^k \mathbb{1}_{\{Y_{kt} = 0\}} \cdot \pi_{1t}^k \mathbb{1}_{\{Y_{kt} > 0\}} \cdot f_s(|y_{kt}| | Y_{kt} \neq 0, \mathcal{F}_{t-1})^{(1 - \mathbb{1}_{\{Y_{kt} = 0\}})}, \end{aligned}$$

where $f_s(|y_{kt}| | Y_{kt} \neq 0, \mathcal{F}_{t-1})$ denotes the conditional density of the size process, i.e. conditional density of an absolute change of Y_{kt} , with support $\mathbb{N} \setminus \{0\}$. To get a parsimoniously specified model, we adopt the simplification of Liesenfeld et al. (2006) that the conditional density of an absolute value of a variable stems from the same distribution irrespective of whether the variable is positive or negative.

The conditional probabilities of the direction process are modelled with the autoregressive conditional multinomial model (ACM) of Russell and Engle (2002) using a logistic link function given by

$$\pi_{jt}^k = \frac{\exp(\Lambda_{jt}^k)}{\sum_{j=-1}^1 \exp(\Lambda_{jt}^k)} \quad (4)$$

where $\Lambda_{0t}^k = 0$, $\forall t$ is the normalizing constraint. The resulting vector of log-odds ratios $\Lambda_t^k \equiv (\Lambda_{-1t}^k, \Lambda_{1t}^k)' = (\ln[\pi_{-1t}^k/\pi_{0t}^k], \ln[\pi_{1t}^k/\pi_{0t}^k])'$ is specified as a multivariate ARMA(1,1) model:

$$\Lambda_t^k = \mu + B_1 \Lambda_{t-1}^k + A_1 \xi_{t-1}^k. \quad (5)$$

μ denotes the vector of constants, and B_1 and A_1 denote 2×2 coefficient matrices. In the empirical application, we put the following symmetry restrictions $\mu_1 = \mu_2$, as well as $b_{11}^{(1)} = b_{22}^{(1)}$ and $b_{21}^{(1)} = b_{12}^{(1)}$ on the B_1 matrix to obtain a parsimonious model specification. The innovation vector of the ARMA model is specified as a martingale difference sequence in the following way:

$$\xi_{jt}^k \equiv (\xi_{-1t}^k, \xi_{1t}^k)', \quad \text{where} \quad \xi_{jt}^k \equiv \frac{x_{jt}^k - \pi_{jt}^k}{\sqrt{\pi_{jt}^k(1 - \pi_{jt}^k)}}, \quad j \in \{-1, 1\}, \quad (6)$$

and

$$x_t^k \equiv (x_{-1t}^k, x_{1t}^k)' = \begin{cases} (1, 0)' & \text{if } Y_{kt} < 0 \\ (0, 0)' & \text{if } Y_{kt} = 0 \\ (0, 1)' & \text{if } Y_{kt} > 0, \end{cases} \quad (7)$$

denotes the state vector, whether Y_{kt} decreases, stays equal or increases at time t . Thus, ξ_t^k represents the standardized state vector x_t^k .

The conditional density of the size process is modelled with an at-zero-truncated Negative Binomial (NegBin) distribution:

$$\begin{aligned} f_s(|y_{kt}| | Y_{kt} \neq 0, \mathcal{F}_{t-1}) \\ \equiv \frac{\Gamma(\kappa^k + |y_{kt}|)}{\Gamma(\kappa^k) \Gamma(|y_{kt}| + 1)} \left(\left[\frac{\kappa^k + \omega_t^k}{\kappa^k} \right]^{\kappa^k} - 1 \right)^{-1} \left(\frac{\omega_t^k}{\omega_t^k + \kappa^k} \right)^{|y_{kt}|}, \end{aligned} \quad (8)$$

where $|y_{kt}| \in \mathbb{N} \setminus \{0\}$, $\kappa^k > 0$ denotes the dispersion parameter and scaling parameter ω_t^k is parameterized using the exponential link function with a generalized autoregressive moving average model (GLARMA(p, q)) of Shephard (1995) in the following way:

$$\ln \omega_t^k = \delta \tilde{D}_t + \tilde{\lambda}_t^k \quad \text{with} \quad \tilde{\lambda}_t^k = \tilde{\mu} + S^k(\nu, \tau, K) + \beta_1 \tilde{\lambda}_{t-1}^k + \alpha_1 \tilde{\xi}_{t-1}^k.$$

where $\tilde{D}_t \in \{-1, 1\}$ indicates a negative or positive value of Y_{kt} at time t with the corresponding coefficient denoted by δ . $\tilde{\mu}$ denotes the constant term. β_1 as well as α_1 denote coefficients and $\tilde{\xi}_t^k$ being constructed as

$$\tilde{\xi}_t^k \equiv \frac{|Y_{kt}| - \mathbb{E}(|Y_{kt}| | Y_{kt} \neq 0, \mathcal{F}_{t-1})}{\mathbb{V}(|Y_{kt}| | Y_{kt} \neq 0, \mathcal{F}_{t-1})^{1/2}},$$

is the innovation term that drives the GLARMA model in λ_t^k . The conditional moments of the at-zero-truncated NegBin distribution are given by

$$\begin{aligned} \mathbb{E}(|Y_{kt}| | Y_{kt} \neq 0, \mathcal{F}_{t-1}) &= \frac{\omega_t^k}{1 - \vartheta_t^k}, \\ \mathbb{V}(|Y_{kt}| | Y_{kt} \neq 0, \mathcal{F}_{t-1}) &= \frac{\omega_t^k}{1 - \vartheta_t^k} - \left(\frac{\omega_t^k}{(1 - \vartheta_t^k)} \right)^2 \left(\vartheta_t^k - \frac{1 - \vartheta_t^k}{\kappa^k} \right), \end{aligned}$$

where ϑ_t^k is given by $\vartheta_t^k = [\kappa^k / (\kappa^k + \omega_t^k)]^\kappa$.

$$S^k(\nu, \tau, K) \equiv \nu_0 \tau + \sum_{l=1}^K \nu_{2l-1} \sin(2\pi(2l-1)\tau) + \nu_{2l} \cos(2\pi(2l)\tau) \quad (9)$$

is a Fourier flexible form used to capture diurnal seasonality, where τ is the intraday time standardized to $[0, 1]$ and ν is a $2K + 1$ dimensional parameter vector.

3 Bivariate modelling of exchange rate changes

Data description

We apply our model to 1 min mid quote changes of the EUR/GBP and the EUR/USD exchange rates. The data has been provided by Olsen Financial Technologies and contains quotes from the electronic foreign exchange interbank market. The period under study spreads between October 6 (Monday), 2003, 0:01 EST, and October 10 (Friday), 2003, 17:00 EST, resulting in 6,780 observations for both time series. The sampling frequency of 1 min is, on the one side, sufficiently high to maintain the discrete nature of the data, whereas on the other side, it is low enough to preserve a significant correlation between the two marginal processes.

The histograms of the two marginal processes are presented in the Fig. 2. Both distributions reveal a fairly large support between -20 and 20 ticks for the EUR/GBP and between -30 and 30 ticks for the more volatile EUR/USD exchange rate. Thus, the discreteness of the quote changes combined with a high number of zero quote movements (about 13% for the EUR/GBP and about 7.5% for the EUR/USD) justifies the ICH-model approach of Liesenfeld et al. (2006).

We associate Y_{1t} and Y_{2t} with the changes of the EUR/GBP and the EUR/USD currency pairs, respectively, and present in Figs. 3 and 4 the dynamic features of these variables in the form of the multivariate autocorrelograms of the vectors x_t^1 and x_t^2 , which indicate the change of the direction of the EUR/GBP and the EUR/USD exchange rates, as defined in Eq. (7).

We observe that there is a certain dynamic pattern, which should be explained by the ACM part of the ICH model. As indicated by the negative first-order autocorrelation and the positive first-order cross correlation coefficients, the probability of an upward movement of each exchange rate following a downward movement is significantly more probable than two subsequent negative or positive movements. In Fig. 5 the autocorrelograms for the absolute value of the non-zero exchange rate changes are presented. The high degree of persistence characterizing the processes should be explained by the GLARMA part of the ICH model.

The interdependence between the two marginal processes can be seen from Fig. 6, where we plotted the multivariate autocorrelogram of Y_{1t} and Y_{2t} . The two marginal processes are positively correlated, with the correlation coefficient of about 0.35.

Fig. 2 Histograms of the tick changes of the EUR/GBP and the EUR/USD exchange rates.

Fig. 3 Multivariate autocorrelation function for the EUR/GBP mid quote direction. *Upper left panel:* $\text{corr}(x_{-1,t}^1, x_{-1,t-l}^1)$; *upper right panel:* $\text{corr}(x_{-1,t}^1, x_{1,t-l}^1)$; *lower left panel:* $\text{corr}(x_{-1,t-l}^1, x_{1,t}^1)$ and *lower right panel:* $\text{corr}(x_{1,t}^1, x_{1,t-l}^1)$. The dashed lines mark the approximate 99% confidence interval $\pm 2.58/\sqrt{T}$.

Fig. 4 Multivariate autocorrelation function for the EUR/USD mid quote direction. *Upper left panel:* $\text{corr}(x_{-1,t}^2, x_{-1,t-l}^2)$; *upper right panel:* $\text{corr}(x_{-1,t}^2, x_{1,t-l}^2)$; *lower left panel:* $\text{corr}(x_{-1,t-l}^2, x_{1,t}^2)$ and *lower right panel:* $\text{corr}(x_{1,t}^2, x_{1,t-l}^2)$. The dashed lines mark the approximate 99% confidence interval $\pm 2.58/\sqrt{T}$.

Fig. 5 Autocorrelation function of the non-zero absolute EUR/GBP and EUR/USD mid quote changes. The *dashed line* marks the approximate 99% confidence interval $\pm 2.58/\sqrt{\tilde{T}}$, where $\tilde{T}$ is the number of non-zero quote changes.

Fig. 6 Multivariate autocorrelation function for the EUR/GBP and EUR/USD mid quote changes. *Upper left panel*: $\text{corr}(Y_{1t}, Y_{1t-1})$; *upper right panel*: $\text{corr}(Y_{1t}, Y_{2t-1})$; *lower left panel*: $\text{corr}(Y_{1t-1}, Y_{2t})$ and *lower right panel*: $\text{corr}(Y_{2t}, Y_{2t-1})$. The *dashed lines* mark the approximate 99% confidence interval $\pm 2.58/\sqrt{\tilde{T}}$.

Bivariate model specification

The copula concept allows one to model the bivariate density without forcing the direction of the dependence upon the data generating process. We choose the standard Gaussian copula function since its single dependency parameter can easily be estimated and it allows for a straightforward sampling algorithm of variables from the bivariate conditional density, which is necessary to assess the goodness-of-fit of our specification. The Gaussian copula is given by

$$\mathbf{C}(u_t, v_t; \rho) = \int_{-\infty}^{\Phi^{-1}(u_t)} \int_{-\infty}^{\Phi^{-1}(v_t)} \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(\frac{2\rho uv - u^2 - v^2}{2(1-\rho^2)}\right) dudv, \quad (10)$$

where $u_t = F(y_{1t}|\mathcal{F}_{t-1})$, $v_t = F(y_{2t}|\mathcal{F}_{t-1})$ and ρ denotes the time-invariant parameter of the Gaussian copula, which is the correlation between $\Phi^{-1}(u_t)$ and $\Phi^{-1}(v_t)$. Since ρ is chosen to be fixed over time $\mathbf{C}(F(y_{1t}|\mathcal{F}_{t-1}), F(y_{2t}|\mathcal{F}_{t-1})|\mathcal{F}_{t-1}) = \mathbf{C}(F(y_{1t}|\mathcal{F}_{t-1}), F(y_{2t}|\mathcal{F}_{t-1}))$ and the conditional bivariate density of Y_{1t} and Y_{2t} can be inferred from Eq. (3) as

$$\begin{aligned} f(y_{1t}, y_{2t}|\mathcal{F}_{t-1}) &= \mathbf{C}(F(y_{1t}|\mathcal{F}_{t-1}), F(y_{2t}|\mathcal{F}_{t-1})) \\ &\quad - \mathbf{C}(F(y_{1t} - 1|\mathcal{F}_{t-1}), F(y_{2t}|\mathcal{F}_{t-1})) \\ &\quad - \mathbf{C}(F(y_{1t}|\mathcal{F}_{t-1}), F(y_{2t} - 1|\mathcal{F}_{t-1})) \\ &\quad + \mathbf{C}(F(y_{1t} - 1|\mathcal{F}_{t-1}), F(y_{2t} - 1|\mathcal{F}_{t-1})). \end{aligned} \quad (11)$$

The cumulative distribution function $F(y_{1t}|\mathcal{F}_{t-1})$ (and analogously $F(y_{2t}|\mathcal{F}_{t-1})$) can be written as

$$\begin{aligned} F(y_{1t}|\mathcal{F}_{t-1}) &= \sum_{k=-50}^{y_{1t}} \pi_{-1t}^1 \mathbb{1}_{\{k<0\}} \cdot \pi_{0t}^1 \mathbb{1}_{\{k=0\}} \cdot \pi_{1t}^1 \mathbb{1}_{\{k>0\}} \cdot f_s(|k| | k \neq 0, \mathcal{F}_{t-1})^{(1-\mathbb{1}_{\{k=0\}})} \end{aligned}$$

where we set the lower bound of the summation to -50 and where the probabilities of the downward, zero and upward movement of the exchange rate are specified with the logistic link function, as shown in Eq. (4), and the density for the absolute value of the change is specified along the conditional NegBin distribution, as presented in Eq. (8).

Estimation and simulation results

Our estimates are obtained by a one-step maximum likelihood estimation, whereas the log-likelihood function is taken as a logarithm of the bivariate density presented in the Eq. (11). Estimation results for the ACM part of ICH model are presented in Table 1 and for the GLARMA part of the ICH model in Table 2. The estimate of the dependency parameter $\hat{\rho}$ for the Gaussian copula equals to 0.3588 with standard

Table 1 ML estimates of the ACM-ARMA part of ICH model. Multivariate Ljung-Box statistic for L lags, $Q(L)$, is computed as $Q(L) = n \sum_{\ell=1}^L \text{tr}[\Gamma(\ell)' \Gamma(0)^{-1} \Gamma(\ell) \Gamma(0)^{-1}]$, where $\Gamma(\ell) = \sum_{i=\ell+1}^n v_i v_{i-\ell}' / (n - \ell - 1)$. Under the null hypothesis, $Q(L)$ is asymptotically χ^2 -distributed with degrees of freedom equal to the difference between 4 times L and the number of parameters to be estimated

	EUR/GBP		EUR/USD	
Parameter	Estimate	Standard deviation	Estimate	Standard deviation
μ	0.0639	0.0177	0.0837	0.0222
$b_{(11)}^{(1)}$	0.6583	0.0856	0.4518	0.0426
$b_{(12)}^{(1)}$	0.2910	0.0540	0.5054	0.0635
$a_{11}^{(1)}$	0.1269	0.0324	0.2535	0.0465
$a_{12}^{(1)}$	0.2059	0.0323	0.3739	0.0472
$a_{21}^{(1)}$	0.2009	0.0271	0.3350	0.0466
$a_{22}^{(1)}$	0.0921	0.0312	0.2586	0.0477
Resid. mean	$(-0.003, 0.002)$		$(0.003, 0.009)$	
Resid. variance	$\begin{pmatrix} 0.655 & 0.803 \\ 0.803 & 2.631 \end{pmatrix}$		$\begin{pmatrix} 1.413 & 2.306 \\ 2.306 & 4.721 \end{pmatrix}$	
Resid. $Q(20)$	72.359 (0.532)		89.054 (0.111)	
Resid. $Q(30)$	102.246 (0.777)		122.068 (0.285)	
Log-lik.	-6.2125			

Table 2 ML estimates of the GLARMA part of ICH model

Parameter	EUR/GBP		EUR/USD	
	Estimate	Standard deviation	Estimate	Standard deviation
$\kappa^{0.5}$	0.7862	0.0192	0.7952	0.0130
$\tilde{\mu}$	0.3363	0.0438	0.7179	0.0814
β_1	0.6567	0.0335	0.6085	0.0428
α_1	0.1675	0.0100	0.1455	0.0097
ν_0	0.0981	0.0633	−0.0396	0.0510
ν_1	−0.0712	0.0117	−0.0430	0.0100
ν_2	−0.0060	0.0091	0.0388	0.0099
ν_3	−0.0501	0.0105	−0.0258	0.0093
ν_4	0.0852	0.0238	0.0954	0.0215
ν_5	−0.0100	0.0129	−0.0234	0.0120
ν_6	0.0449	0.0115	0.0297	0.0108
Resid. mean	0.013		0.007	
Resid. variance	1.001		1.025	
Resid. $Q(20)$	26.408 (0.067)		42.332 (0.001)	
Resid. $Q(30)$	64.997 (0.000)		87.641 (0.000)	
Log-lik.	−6.2125			

deviation 0.0099, representing a strong positive correlation between the modelled marginal processes.

Regarding the estimates for the ACM submodel, we observe a significant persistency pattern ($\hat{B}_1$ matrix) of the direction processes and we can conclude, that if the probability of an exchange rate change has been high in the previous period, it is also supposed to be considerably high in the next period. Moreover, the obtained relations $a_{11}^{(1)} < a_{12}^{(1)}$ and $a_{21}^{(1)} > a_{22}^{(1)}$ between the innovation coefficients suggest the existence of some bounce or mean-reverting pattern in the evolution of the exchange rate process. The parsimonious dynamic specification seems to describe the dynamic structure very well, as the multivariate Ljung-Box statistics for the standardized residuals of the ACM model do not differ significantly from zero.

Regarding the estimation results for the GLARMA part of the ICH model, we observe that the values of the dispersion parameters $\kappa^{k-0.5}$ are significantly different from zero, allowing the rejection of the null hypothesis of at-zero-truncated Poisson distributions in favor of at-zero-truncated NegBin ones. The diagnostics statistics of the GLARMA submodel are quite satisfying. Although some Ljung-Box statistics (Q) for the standardized residuals still remain significantly different from zero, a large part of the autocorrelation structure of the size processes has been explained by the simple GLARMA(1,1) specification.

Jointly significant coefficients of the seasonal components $S(\nu, \tau, K)$ for $K = 3$ indicate that there exist diurnal seasonality patterns, which are plotted in Fig. 7, in the absolute changes of the exchange rates. We observe for every minute of the day that the mean of the non-zero absolute tick changes of the USD against the EUR is considerably higher than the mean for the GBP against the EUR. It confirms the results of the descriptive study presented previously, as the support of the EUR/USD distribution is more dispersed and the exchange rate is more volatile. The shapes of the diurnal seasonality functions for both exchange rates are quite similar. They evidence the existence of at least two very active trading periods, about 3.00 EST and 10.00 EST, which corresponds to the main trading periods of the European and the American Foreign Exchange market, respectively.

In order to verify the goodness-of-fit of our model in a more elaborate way, we simulate the conditional density of the bivariate process at every

Fig. 7 Estimated diurnal seasonality function of the non-zero absolute EUR/GBP and EUR/USD tick changes. The *dashed line* mark the approximate 99% confidence interval.

point t , $t = 1, \dots, T$. Such an approach enables us to verify whether the proposed density specification is able to explain the whole conditional joint density of the underlying data generating process. Relying on the simulated distributions at every time point available, we can easily address this point applying the modified version of the Diebold et al. (1998) density forecasting test for discrete data of Liesenfeld et al. (2006).

Moreover, we are able to compare the residuals of both marginal processes. We use here the standard sampling method proposed for Gaussian copula functions, which can be summarized as:

For every t :

- Compute the Cholesky decomposition $\hat{A}$ (2×2) of estimated correlation matrix $\hat{R}$, where $\hat{R} = \begin{pmatrix} 1 & \hat{\rho} \\ \hat{\rho} & 1 \end{pmatrix}$
- Simulate $x_t = (x_{1t}, x_{2t})'$ from a bivariate standard normal distribution

Fig. 8 Histogram of simulated tick changes of exchange rates.

Fig. 9 Bivariate histogram of the simulated tick changes of the EUR/GBP and the EUR/USD exchange rates.

- Set $\hat{z}_t = \hat{A}x_t$
- Set $\hat{u}_{1t} = \Phi(\hat{z}_{1t})$ and $\hat{u}_{2t} = \Phi(\hat{z}_{2t})$ where Φ denotes the univariate standard normal distribution function
- Set $\hat{Y}_{1t} = \hat{F}_1^{-1}(\hat{u}_{1t}|\mathcal{F}_{t-1})$ and $\hat{Y}_{2t} = \hat{F}_2^{-1}(\hat{u}_{2t}|\mathcal{F}_{t-1})$ where $\hat{F}_1$ and $\hat{F}_2$ denote the estimated marginal cumulative distribution functions of the EUR/GBP and the EUR/USD changes, respectively

Figure 8 contains the plots of the unconditional histograms of the simulated marginal processes. Their shapes seem to agree with those of the raw data series already presented in Fig. 2.

Fig. 10 Bivariate histogram of the positive differences between the empirical and the simulated bivariate histogram of the EUR/GBP and the EUR/USD exchange rate changes.

Fig. 11 Bivariate histogram of the absolute values of the negative differences between the empirical and the simulated bivariate histogram of the the EUR/GBP and the EUR/USD exchange rate changes.

The unconditional bivariate histogram of the simulated time series is presented in Fig. 9. Although the positive dependence between the marginal processes is reflected, the shape of the histogram does not correspond to the empirical one in full (see Fig. 1). In particular the frequency of the outcome $(0, 0)$ has been considerably underestimated. We compute the differences between the histograms of the empirical and the simulated data to infer in which points (i, j) the observed and the estimated probabilities disagree. To assess these differences graphically, we plotted in Fig. 10 only positive differences and in Fig. 11 only absolute negative differences. Besides the outcome probability of $(0, 0)$, the probabilities for points (i, j) concentrated around $(0, 0)$ are a little bit underestimated (positive differences in Fig. 10) as well, and the probabilities for points (i, j) which are a little further away from $(0, 0)$ are a little overestimated (negative differences in Fig. 11). Thus, we conclude that we underestimate the kurtosis of the empirical distribution. The real data is much more concentrated in the outcome $(0, 0)$, as well as evidencing much fatter tails. There is a clear signal for a tail dependency in the data generating process, as the extreme positive or negative movements of the exchange rates take place much more often than could be explained by a standard Gaussian copula function (see Fig. 10).

Additionally, we can address the goodness-of-fit of the conditional bivariate density by considering the bivariate autocorrelation function of the residual series $\hat{\varepsilon}_t = (\hat{\varepsilon}_t^1, \hat{\varepsilon}_t^2)'$ depicted in Fig. 14 and the quantile–quantile (QQ) plots of the modified density forecast test variables for the implied marginal processes in Fig. 12. We have mapped these modified density forecast test variables into a standard normal distribution, so that under the correct model specification, these normalized variables should be i.i.d. standard normally distributed. Figure 13 plots the autocorrelation functions of these normalized density forecast variables. Both plots indicate that the processes are almost uncorrelated. The deviation from normality, especially for the EUR/USD exchange rate changes and in the upper tail of the normalized density forecast variables indicated by the QQ-plots, reveals that our specification has difficulties to characterize extreme exchange rate changes appropriately.

The bivariate autocorrelation function of the residual series (Fig. 14) shows significant cross-correlations at lag 1. Although, we manage to explain a large

Fig. 12 QQ plot of the normalized density forecast variables.

Fig. 13 Autocorrelation function of the normalized density forecast variables.

Fig. 14 Multivariate autocorrelation function for residuals of the EUR/GBP and the EUR/USD tick changes. *Upper left panel:* $\text{corr}(\varepsilon_t^1, \varepsilon_{t-1}^1)$; *upper right panel:* $\text{corr}(\varepsilon_t^1, \varepsilon_{t-1}^2)$; *lower left panel:* $\text{corr}(\varepsilon_{t-1}^2, \varepsilon_t^1)$ and *lower right panel:* $\text{corr}(\varepsilon_t^2, \varepsilon_{t-1}^2)$. The dashed lines mark the approximate 99% confidence interval $\pm 2.58/\sqrt{T}$.

part of the correlation structure of the processes for the exchange rate changes (compare Fig. 6), there is some room to improve the model specification. These results are also emphasized by the multivariate Ljung-Box statistics for the bid and ask change process and its residuals in Table 3.

Table 3 Multivariate Ljung-Box statistic for the residuals of the simulated bivariate process

	Real exchange rate changes	Residuals
$Q(20)$	473.6279	159.8679
$Q(30)$	521.8745	207.2940
$Q(50)$	588.0909	296.7522

4 Conclusion

In this paper we propose an approach that is capable of modelling complex multivariate processes for discrete random variables. Combining the approach by Cameron et al. (2004) for copulas of discrete random variables with the ICH model by Liesenfeld et al. (2006), we model the joint process for two integer count variables.

As an illustration of the explanatory power of our approach we estimate the joint distribution of the EUR/GBP and the EUR/USD exchange rate changes at the 1 min level. Even without detailed specification search, our model describes the exchange rate dynamics fairly well. Moreover, the marginal distributions which are characterized by inflated outcomes are also estimated satisfactorily.

In order to pick up the obvious excess kurtosis in the joint empirical distribution, we have tried out more flexible parametric alternatives to the Gaussian copula, such as the t-student copula, which allows for symmetric lower and upper tail dependency and an excessive concentration in (0,0) and the symmetrized Joe–Clayton copula, which has a quite parsimonious functional form and allows for asymmetric tail dependence. Although both specifications improve the goodness-of-fit of our model in some aspects the application of the t-student copula or the symmetrized Joe–Clayton copula has been by no means clearly superior to the simple Gaussian copula, so that we conclude that simply applying more flexible copula functions is not the proper remedy to capture the large excess kurtosis. An obvious alternative path of future research is to keep the Gaussian copula and to inflate the outcome (0,0) along the lines of zero inflated count data models.

Last but not least, the potential merits of the approach should be checked in the light of real world applications such as the measurement of multivariate conditional volatilities and the quantification of liquidity or value-at-risk applications. Obviously, our approach can easily be extended to the most general case of mixed multivariate distributions for continuous and discrete random variables. For financial market research at the high frequency level, such an extension is attractive for the joint analysis of several marks of the trading process (volumes, price and volume durations, discrete quote changes, etc.).

Acknowledgements The work is supported by the European Community's Human Potential Program under contract HPRN-CT-2002-00232, Microstructure of Financial Markets in Europe; and by the Fritz Thyssen Foundation through the project 'Dealer-Behavior and Price-Dynamics on the Foreign Exchange Market'. We thank Richard Olsen and Olsen Financial Technologies for providing us with the data.

References

- Amilon H (2003) GARCH estimation and discrete stock prices: an application to low-priced Australian stocks. *Econ Lett* 81(2):215–222
- Andersen TG, Bollerslev T, Diebold FX, Labys P (1999) (Understanding, optimizing, using and forecasting) realized volatility and correlation, New York University, Leonard N. Stern School Finance Department Working Paper, No. 99–061
- Antoniou A, Vorlow CE (2005) Price clustering and discreteness: is there chaos behind the noise? *Physica A* 348:389–403
- Ball C (1988) Estimation bias induced by discrete security prices. *J Finance* 43:841–865
- Brock WA, Dechert WD, Scheinkman JA, LeBaron B (1996) A test for independence based on the correlation dimension. *Econ Rev* 15(3):197–235
- Cameron C, Li T, Trivedi P, Zimmer D (2004) Modelling the differences in counted outcomes using bivariate copula models with application to mismeasured counts. *Econ J* 7:566–584
- Crack TF, Ledoit O (1996) Robust structure without predictability: the “compass rose” pattern of the stock market. *J Finance* 51(2):751–762
- Denuit M, Lambert P (2005) Constraints on concordance measures in bivariate discrete data. *J Multivariate Anal* 93:40–57
- Diebold FX, Gunther TA, Tay AS (1998) Evaluating density forecasts, with applications to financial risk management. *Int Econ Rev* 39:863–883
- Fang Y (2002) The compass rose and random walk tests. *Comput Stat Data Anal* 39:299–310
- Gleason KC, Lee CI, Mathur I (2000) An explanation for the compass rose pattern. *Econ Lett* 68(2):127–133
- Hansen PR, Lunde A (2006) Realized variance and market microstructure noise. *J Bus Econ Stat* 24:127–218
- Heinen A, Rengifo E (2003) Multivariate autoregressive modelling of time series count data using copulas, Center for Operations Research and Econometrics, Catholique University of Luvain
- Huang RD, Stoll HR (1994) Market microstructure and stock return predictions. *Rev Financ Stud* 7(1):179–213
- Johnson N, Kotz S, Balakrishnan N (1997) Discrete multivariate distributions. Wiley, New York
- Kocherlakota S, Kocherlakota K (1992) Bivariate discrete distributions. Dekker, New York
- Krämer W, Runde R (1997) Chaos and the compass rose. *Econ Lett* 54(2):113–118
- Lee CI, Gleason KC, Mathur I (1999) A comprehensive examination of the compass rose pattern in futures markets. *J Futures Mark* 19(5):541–564
- Liesenfeld R, Nolte I, Pohlmeier W (2006) Modelling financial transaction price movements: a dynamic integer count data model. *Empir Econ* 30:795–825
- Meester S, MacKay J (1994) A parametric model for cluster correlated categorical data. *Biometrics* 50:954–963
- Oomen RCA (2005) Properties of bias-corrected realized variance under alternative sampling schemes. *J Financ Econ* 3:555–577
- Patton A (2001) Modelling time-varying exchange rate dependence using the conditional copula. Discussion Paper, UCSD Department of Economics
- Russell JR, Engle RF (2002) Econometric analysis of discrete-valued irregularly-spaced financial transactions data. University of California, San Diego, Revised Version of Discussion Paper, No. 98–10
- Shephard N (1995) Generalized linear autoregressions. Working Paper, Nuffield College, Oxford
- Sklar A (1959) Fonctions de répartition à n dimensions et leurs marges. *Public Institute of Statistics at the University of Paris* 8:229–231
- Stevens W (1950) Fiducial limits of the parameter of a discontinuous distribution. *Biometrika* 37:117–129
- Szipo GG (1998) Tick size, the compass rose and market nanostructure. *J Bank Finance* 22(12):1559–1569
- Vorlow CE (2004) Stock price clustering and discreteness: the “compass rose” and predictability. Working Paper, University of Durham